arXiv:1608.01074v1 [cs.CV] 3 Aug 2016

FPGA system for real-time computational extended depth of field imaging using phase aperture coding

Tal Remez, Or Litany, Shachar Yoseff, Harel Haim, and Alex Bronstein

Tel-Aviv University
Electrical Engineering Department, Tel Aviv University, Tel Aviv, Israel

Abstract. We present a proof-of-concept end-to-end system for computational extended depth of field (EDOF) imaging. The acquisition is performed through a phase-coded aperture implemented by placing a thin wavelength-dependent optical mask inside the pupil of a conventional camera lens, as a result of which, each color channel is focused at a different depth. The reconstruction process receives the raw Bayer image as the input, and performs blind estimation of the output color image in focus at an extended range of depths using a patch-wise sparse prior. We present a fast non-iterative reconstruction algorithm operating with constant latency in fixed-point arithmetics and achieving real-time performance in a prototype FPGA implementation. The output of the system, on simulated and real-life scenes, is qualitatively and quantitatively better than the result of clear-aperture imaging followed by state-of-the-art blind deblurring.

Keywords: EDOF, sparse coding, FPGA, real-time.

1 Introduction

The increasing popularity of phone cameras combined with the affordability of high-quality CMOS sensors have transformed digital photography into an integral part of the way we communicate. While being the principal technology driver for miniature cameras, the smartphone market is also one of the most challenging arenas for digital imaging. The quality of a digital image is determined by the quality of the optical system and the image sensor. With the increase in pixel number and density, imaging system resolution is now mostly bound by optical system limitations. Since form factor challenges imposed by smart phone designs make it very difficult to improve the image quality via standard optical solutions, many R&D activities in the field in recent years have shifted to the domain of computational imaging. The need to acquire high-quality images and videos of moving low-light scenes through a miniature lens and the fact that traditional autofocusing mechanisms are costly, slow, and unreliable in low-light conditions render acute the tradeoff between the aperture size (F#) and the depth of field (DOF) of the optical system.

Imaging with limited DOF brings forth the challenge of restoration of out-of-focus (OOF) images – a notoriously ill-posed problem due to information loss in the process. There exists a wealth of literature dedicated to purely computational approaches to image deblurring and deconvolution [1,2,3,4]. The dominant models, increasingly popular in recent years, are flavors of sparse and redundant representations [5,6], which

have been proved as a powerful tool for image processing, compression, and analysis. It is now well-established that small patches from a natural image can be represented as a linear combination of only a few atoms in an appropriately constructed over-complete (redundant) dictionary. This constitutes a powerful prior that has been successfully employed to regularize numerous otherwise ill-posed image processing and restoration tasks [7,8].

Existing image deblurring techniques assume that the blurring kernel is known [5] or use an expectation maximization-type of approach to blindly estimate the kernel from the data, either parametrically when the point spread function (PSF) of the system is well-characterized and only the scene depth is unknown, or non-parametrically making less assumptions about the optics [1,6,2,3]. Using conventional optics, information about sharp image features is irreversibly lost, posing inherent limitations on any computational technique.

These limitations can be overcome by manipulating the image acquisition process. Some well-known methods produced a constant PSF kernel for all depths of field using a wave-front coding phase mask [9,10]. Similar methods for depth-invariant PSF utilize focal sweep [11] or use uncorrected lens as a type of spectral focal sweep [12]. Having a single known PSF, the image can be recovered using non-blind deconvolution methods. These techniques reduce image contrast and are very sensitive to noise and therefore not suitable for high quality image applications.

Other related techniques use an amplitude coded mask [13] or a color-dependent ring mask [14] such that objects at different depth will exhibit a distinctive spatial structure which improves the image restoration process. The main drawback of those methods is the fact that the light efficiency is only about 50% in [13] and 60% in [14], making them unsuitable for low light conditions.

A recent trend in the literature [15] proposes the use of wavelength-dependent phase aperture coding, whereby one gets different responses for the red, green and blue channels, resulting in the simultaneous acquisition of three perfectly registered images, each with a different out-of-focus characteristic. Under the assumption that sharp edges and other high-frequency features characterizing a well-focused image are present in all color channels, the joint analysis of the three channels acquired through a phase mask located in pupil, allows to extend the system DOF without significant prior knowledge about the blurring kernel. This type of mask can be easily assembled in most conventional systems with minimal modification and provides a light efficiency of up to 99% using a coated glass plate.

Contributions presented in this work are threefold:

1. We present a learning-based algorithm for computational EDOF using phase aperture coding.
2. We demonstrate empirically the superiority of our algorithm over previous methods in reconstruction quality at a significantly lower runtime.
3. We introduce a full end-to-end hardware system including a FPGA implementation of our algorithm.

The rest of the paper is organized as follows: Section 2 details the image formation model in phase-coded aperture acquisition. Section 3 formulates the reconstruction optimization problem, and Section 4 details the fast algorithm for its approximate

solution. Section 5 describes the FPGA system hosting the reconstruction process. Experimental evaluation of the system on synthetic and real scenes, and its comparison to clear-aperture imaging following by standard deconvolution is presented in Section 6. A brief discussion in Section 7 concludes the paper.

2 Phase-Coded Aperture Imaging

An imaging system acquiring an out-of-focus object suffers from aberrations, in particular blur, that degrade the image quality resulting in low contrast, loss of sharpness and even loss of information. In digital systems, image blur (main OOF effect) is not observable as long as the image size of a point source in the object plane is smaller than the pixel size in the detector plane. The OOF error is analytically a wave-front error, expressed as a quadratic phase error in the pupil plane [16]. In case of a circular aperture with radius *R*, we define the defocus parameter as

$$\begin{aligned}\psi &= \frac{\pi R^2}{\lambda}\left(\frac{1}{z_{\mathrm{o}}}+\frac{1}{z_{\mathrm{img}}}-\frac{1}{f}\right)=\frac{\pi R^2}{\lambda}\left(\frac{1}{z_{\mathrm{img}}}-\frac{1}{z_{\mathrm{i}}}\right)\\ &=\frac{\pi R^2}{\lambda}\left(\frac{1}{z_{\mathrm{o}}}-\frac{1}{z_{\mathrm{n}}}\right),\end{aligned} \tag{1}$$

where z_{img} is the sensor plane location of an object in the nominal position z_{n}, z_{i} is the ideal image plane for an object located at z_{o} , and λ is the optical wavelength. The defocus parameter ψ measures the maximum phase error at the aperture edge. For $\psi > 1$, the image will experience contrast loss, while for $\psi > 4$ it will experience information loss and even reversal of contrast at some frequencies, as can be observed in Figure 1(a). For a circular clear aperture, the cut-off spatial frequency due to diffraction limit is given by $f_{\mathrm{c}} = \frac{2R}{\lambda z_{\mathrm{i}}}$. Increasing the aperture size R is hence a double-edged sword: it results in higher f_{c} and hence improved contrast and higher signal-to-noise ratio (SNR) due to the increase in the amount of collected light; however, this comes at the cost of a steeper increase in ψ as a function of the deviation of the object depth z_{o} from the nominal depth z_{n}, hence, lower DOF. comes at the cost of diminishing the DOF.

Radially symmetric binary optical phase masks have been proposed for overcoming limitations set by OOF imaging [17]. Milgrom *et al.* [18] proposed the use of a special RGB phase mask that exhibits significantly different response in the three major color channels R, G and B. It has been shown that each channel provides best performance for different depth regions, so that the three channels jointly provide an extended DOF. Based on Milgrom's mask, a similar mask was designed in [19] specifically to further increase the diversity between the color channels such that each depth region will be as unique as possible. This effect has been the key element of developing this new type of computational camera. Figure 1(b) shows the diversity between the color channels for different depths (ψ values) when using a clear aperture (dotted plot) and when using the phase mask from [19] (solid plot). The contrast levels are generally somewhat lower when using the phase mask, for the "in focus" condition, but it exhibits higher contrast

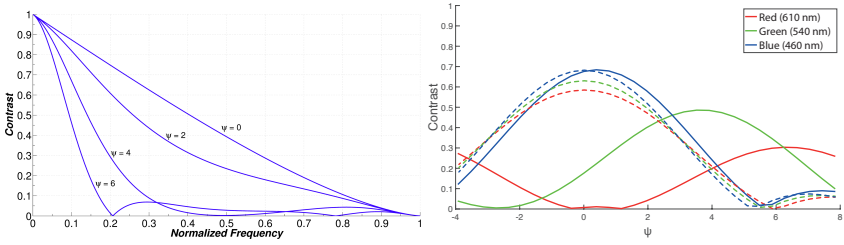

(a) (b)

Fig. 1. Spatial frequency response and color channel separation (a) Optical system response to spatial frequency for different ψ values. (b) Comparison between contrast levels for a single spatial frequency $f_c/4$ as a function of ψ with clear aperture (dotted) and a phase mask (solid).

at larger DOF, providing separation of the response for the three color channels, an effect that will later be used for restoring the contrast levels to their nominal values.

Assuming approximately constant depth and, hence, the same unknown ψ for all pixels in the vicinity of a spatial location $\boldsymbol{\xi}$ in the sensor plane, the image formation model can be described as the following convolution

$$\begin{aligned} y^{\mathrm{R}}(\boldsymbol{\xi}) &= (h_{\psi}^{\mathrm{R}} * x^{\mathrm{R}})(\boldsymbol{\xi}) + \eta^{\mathrm{R}}(\boldsymbol{\xi}) \\ y^{\mathrm{G}}(\boldsymbol{\xi}) &= (h_{\psi}^{\mathrm{G}} * x^{\mathrm{G}})(\boldsymbol{\xi}) + \eta^{\mathrm{G}}(\boldsymbol{\xi}) \\ y^{\mathrm{B}}(\boldsymbol{\xi}) &= (h_{\psi}^{\mathrm{B}} * x^{\mathrm{B}})(\boldsymbol{\xi}) + \eta^{\mathrm{B}}(\boldsymbol{\xi}), \end{aligned}$$

where $x^*, * = R, G, B$, denote the three color channels of the ideal (latent) in-focus image, y^* are the corresponding channels of the OOF image formed in the sensor plane, h_{ψ}^* are the three corresponding depth-dependent blur kernels, and η^* denote additive sensor noise that can be fairly accurately modeled as white Gaussian at least at reasonable illumination intensities. In vector notation combining the three channels, the expression simplifies to

$$\mathbf{y}(\boldsymbol{\xi}) = (\mathbf{h}_{\psi} * \mathbf{x})(\boldsymbol{\xi}) + \boldsymbol{\eta}(\boldsymbol{\xi}), \tag{2}$$

where $\mathbf{x}$ and $\mathbf{y}$ denote the three-channel latent in-focus and the observed OOF images, respectively.

In the majority of conventional image sensors, the acquisition of color information is done through a color filter array (CFA), also known as Bayer filter mosaic, where every pixel captures only one of the red, green or blue channels. A demosaicing process then restores the raw image to a full-resolution color image. With some abuse of notation, CFA can be introduced to our image formation model as

$$y(\boldsymbol{\xi}) = \mathbf{B}(\boldsymbol{\xi})(\mathbf{h}_{\psi} * \mathbf{x})(\boldsymbol{\xi}) + \boldsymbol{\eta}'(\boldsymbol{\xi}), \tag{3}$$

where y denotes the raw Bayer image, and $\mathbf{B}(\boldsymbol{\xi})$ the location-dependent response of the array filters, which up to cross-talk and vignetting artifacts looks like $\mathbf{B} = (1,0,0), (0,1,0)$, and $(0,0,1)$ for the R, G, and B color channels, respectively.

3 EDOF Image Recovery Using a Sparse Prior

Sparse representation was proved to be a strong prior for both non-blind [5] and blind [1,6,2,3] image deblurring. A signal $\mathbf{x} \in \mathbb{R}^n$ is said to admit a sparse representation (or, more accurately, approximation) in an $n \times k$ overcomplete ($k > n$) dictionary $\mathbf{D}$ if one can find a vector $\mathbf{z} \in \mathbb{R}^k$ with only a few non-zero coefficients, such that $\mathbf{x} \approx \mathbf{D}\mathbf{z}$ e.g. in the ℓ_2-sense. The sparse representation pursuit problem can be cast as the minimization problem,

$$\hat{\mathbf{z}} = \arg\min_{\mathbf{z}} \|\mathbf{x} - \mathbf{D}\mathbf{z}\|_2^2 + \mu g(\mathbf{z}), \tag{4}$$

where $g(\mathbf{z})$ is a sparsity-inducing prior whose relative importance is governed by the parameter μ.

While priors such as the ℓ_0 pseudonorm counting the number of non-zeros in the vector $\mathbf{z}$ give rise to computationally intractable optimization problems, there exist efficient greedy approximation techniques including the orthogonal matching pursuit (OMP) [20,21] or convex relaxation techniques replacing the ℓ_0 pseudonorm with the ℓ_1 norm minimization [22]. The dictionary $\mathbf{D}$ can be constructed axiomatically based on image transforms such as DCT or wavelet, or trained (e.g., via k-SVD [23]) using a set of representative images. Here, we adopt the latter approach to construct a dictionary representing 8×8 patches from the image, represented as $n = 192$-dimensional vectors (64 dimensions for each color channel).

Assuming (unrealistically) a known ψ, the forward model (3) described in the previous section can be directly incorporated into the data fitting term, resulting in the following optimization problem:

$$\hat{\mathbf{z}} = \arg\min_{\mathbf{z}} \|\mathbf{y} - \mathbf{B}\mathbf{H}_\psi\mathbf{D}\mathbf{z}\|_2^2 + \mu g(\mathbf{z}), \tag{5}$$

where $\mathbf{y}$ is a 64-dimensional vector representing the patch in the input Bayer image, $\mathbf{D}$ is the $192 \times k$-dimensional dictionary, and $\mathbf{z}$ is the k-dimensional vector of representation coefficients. The 192×192 matrix $\mathbf{H}_\psi$ represents the action of the out-of-focus blur kernel corresponding to ψ to each of the color channels, while 64×192 matrix $\mathbf{B}$ describes the action of the CFA on the patch; in conventional CFA designs it is constant as long as the patches are selected from essentially same locations. The reconstructed patch is obtained as $\hat{\mathbf{x}} = \mathbf{D}\hat{\mathbf{z}}$. Defining the blurred dictionary $\mathbf{D}_\psi = \mathbf{H}_\psi\mathbf{D}$, the latter optimization problem can be rewritten as

$$\hat{\mathbf{z}} = \arg\min_{\mathbf{z}} \|\mathbf{y} - \mathbf{B}\mathbf{D}_\psi\mathbf{z}\|_2^2 + \mu g(\mathbf{z}). \tag{6}$$

Note that the pursuit is performed with respect to the blurred dictionary $\mathbf{D}_\psi$, while the patch reconstruction from the computed coefficients is performed with the clear dictionary $\mathbf{D}$.

In practical applications, the depth and, hence, the out-of-focus parameter ψ in each patch are unknown *a priori*, leading to the blind setting of the above reconstruction problem. Many studies dealt with this setting with limited success. For instance, using different iterative processes, one tries to estimate the blurring kernel so that image restoration can be thereafter achieved. Reconstruction processes usually require high

computational complexity, which limits their use for many real-time applications. Even then, the assumption of a single blur kernel for the entire image is often unrealistic.

In our setup, the phase coding of the aperture contains hints about ψ through the relation between the color channels. To address the blind setting of the problem, we construct a concatenated blurred representation dictionary containing several fixed values of ψ,

$$\mathbf{D}_{\Psi} = (\mathbf{D}_{\psi_1}, \ldots, \mathbf{D}_{\psi_q}),$$

and solving

$$\hat{\mathbf{z}} = \arg\min_{\mathbf{z}} \|\mathbf{y} - \mathbf{B}\mathbf{D}_{\Psi}\mathbf{z}\|_2^2 + \mu g(\mathbf{z}), \tag{7}$$

for $\mathbf{z} \in \mathbb{R}^{qk}$. The reconstruction is performed as before with the clear dictionary concatenated q times, $\hat{\mathbf{x}} = (\mathbf{D}, \ldots, \mathbf{D})\hat{\mathbf{z}}$. In all our experiments, we used $q = 8$ and set $\psi_1 = 1, \ldots, \psi_8 = 8$. While an additional prior asserting that the coefficients are localized in one of the sub-dictionaries can be imposed using standard structured sparsity-promoting priors such as the mixed $\ell_{2,1}$ norm, we defer this obvious extension to future study.

The pursuit process chooses elements from the dictionary that best match the input patch, based on the ℓ_2 distance. Using the RGB phase mask, the blurred dictionary varies strongly for different kernels since the response of the imaging system is very different for each color. An input patch from an imaging system phase-coded aperture will also exhibit different response for each color. The response will be unique for each kernel and therefore the input vector will most likely associate with vectors from the dictionary $\mathbf{D}_{\psi_i}$ experiencing the same blurring process.

Comparing our algorithm with a state-of-art algorithm with available public domain code, provided by Krishnan et al. [2], and our implementation of [15], our process produced superior results when applied to natural images out of the Kodak dataset [24]. We also run the process on texture images (Colored Brodatz Texture database [25]) and observed similar performance, meaning that the concatenated dictionary is likely to describe well almost any natural scene.

4 Fast Image Reconstruction

As was explained in the previous section, the pursuit problem (7) can be posed as a convex optimization problem by choosing $g(\mathbf{z}) = \|\mathbf{z}\|_1$, and solved using proximal algorithms such as the iterative shrinkage thresholding algorithm (ISTA) or its accelerated version (FISTA) [26]. However, these iterative solvers typically require hundreds of iterations to converge, resulting in prohibitive complexity and unpredictable input-dependent latency, which is unacceptable in real-time applications. To overcome this limitation, we follow the approach advocated by [27], in which a small number, T, of ISTA iterations is unrolled into a feed-forward neural network that subsequently undergoes supervised training on typical inputs, as explained in the sequel. A pseudo-code of ISTA is given in Algorithm 1, where we denote $\mathbf{Q} = \frac{1}{L}\mathbf{B}\mathbf{D}_{\Psi}$, $\mathbf{S} = \mathbf{I} - \frac{1}{L}\mathbf{Q}^{\mathrm{T}}\mathbf{Q}$, and $\sigma_{\theta}(x) = \max(|x| - \theta, 0)\mathrm{sign(x)}$ is a two-sided shrinkage function with threshold $\theta = \frac{2\mu}{L}$ applied element-wise. L denotes a scalar larger than the largest eigenvalue of $\mathbf{Q}^{\mathrm{T}}\mathbf{Q}$.

Input: Data $\mathbf{x}$, number of iterations T, shrinkage threshold θ, matrices $\mathbf{D}$, $\mathbf{Q}$, and $\mathbf{S}$
Output: Reconstructed image $\hat{\mathbf{x}}$,
auxiliary variables $\{\mathbf{z}_t\}_{t=1}^{T}$,$\{\mathbf{b}_t\}_{t=1}^{T}$
initialize $\mathbf{z}_1 = \bar{0}$ and $\mathbf{b}_1 = \mathbf{Q}^{\mathrm{T}}\mathbf{x}$
for *t = 1,2,...,T* **do**
  $\mathbf{z}_{t+1} = \sigma_\theta\left(\mathbf{b}_t\right)$
  $\mathbf{b}_{t+1} = \mathbf{b}_t + \mathbf{S}\left(\mathbf{z}_{t+1} - \mathbf{z}_t\right)$
end
$\hat{\mathbf{x}} = \mathbf{D}\mathbf{z}_T$

Algorithm 1: ISTA algorithm.

As can be easily interpreted from the ISTA algorithm, the network comprises three types of layers: An initialization layer (denoted as I), which finds the representation of the input signal in the dictionary; several $(T-2)$ recurrent middle layers (M) performing gradient step followed by shrinkage; and a final layer (F) which translates the resulting dictionary coefficients to the reconstructed signal. All these types of layers can be realized from the single multi-purpose calculator stage shown in Figure 2 (right) that is amenable for hardware implementation. To get the "I" configuration, we set $\mathbf{b}_{\mathrm{in}} = \mathbf{0}$, $\mathbf{A} = -\mathbf{Q}^{\mathrm{T}}$, $\mathbf{c} = \mathbf{0}$ and $\boldsymbol{\theta} = \mathbf{0}$. The "M" configuration layer is fed by the output of the previous calculator, and the matrix $\mathbf{A}$ is set to $\mathbf{S}$. The output is further fed to either another "M" layer or to the "F" layer. Finally, the "F" configuration of the calculator is a reduction into multiplication by the matrix $\mathbf{D}$. This is achieved by setting $\mathbf{b}_{\mathrm{in}} = \mathbf{0}$, and $\mathbf{A} = \mathbf{D}$.

Supervised training of the network is done by initializing the parameters as detailed above, and then adapting them using a stochastic gradient procedure minimizing the reconstruction error $\mathcal{F}$ of the entire network. We use the following empirical loss

$$\mathcal{F} = \frac{1}{N}\sum_{n=1}^{N} f(\mathbf{x}_n^*, \hat{\mathbf{x}}_n), \tag{8}$$

which for a large enough training set, N, approximates the expected value of f with respect to the distribution of the ground truth signals $\mathbf{x}_n^*$. Here, $\hat{\mathbf{x}}_n$ denotes the output of the network, and the loss objective f minimized during the training process is the standard sum of squared differences,

$$f = \frac{1}{2}\|\mathbf{x}_n^* - \hat{\mathbf{x}}_n\|_2^2. \tag{9}$$

Similarly to [27], the output of the network and the derivatives of the loss with respect to the network parameters are calculated using the standard forward and back propagation approach. Practice shows that the training process allows to reduce the number of layers by about two orders of magnitude while achieving a comparable reconstruction quality. A detailed runtime analysis is presented in Section 6.

5 FPGA Image Reconstruction System

To demonstrate that the proposed image reconstruction process is efficient and is amenable to hardware implementation, we built a prototype FPGA system. An FPGA is a programmable chip containing configurable logic blocks and routing resources, therefore offering a fair amount of flexibility previously only possible with software, along with a hardware-like computational speeds and reliability. While being distinct in many aspects from application-specific integrated chips (ASICs), modern FPGAs are the closest approximation of an ASIC one can get without incurring the costs of custom chip manufacturing.

A schematic description of our system is depicted in Figure 2 (left). We used the Xilinx Kintex 7 chip on the KC705 development board chosen mainly because of the availability of video interfaces. As the output, we used the onboard HDMI output phy, while for the input, we added an external HDMI phy board connected to the main board through an FCM connector. The input frames are received by the board through the HDMI interface in raw Bayer format, 16 bits per pixel with the most significant bytes packaged as the Y channel, and the least significant byte packaged as the 4:2:2 color channels. The input is relayed to Write Agent 0 on the FPGA chip that buffers it in the external dynamic memory. The content of the buffer is brought into the chip by Read Agent 0, which reorders the raster scan order into a stream of 8×8 patches with configurable amount of overlap. The patches are then fed into a configurable calculator pipeline implementing the reconstruction algorithm detailed in the previous section. The pipeline comprises three configurable stages, one of which is configured as the initial stage (I), another as a middle stage (M), and yet another as the final stage (F), yielding the flow structure of the form $\mathrm{I} \rightarrow (T-2) \times \mathrm{M} \rightarrow \mathrm{F}$.

The output of the calculator pipeline is produced in 4:2:2 YCbCr format comprising 64 luma values at 16 bits per pixel, and additional $32 + 32 = 64$ chroma values at 8 bits per pixel. The luma component undergoes gamma conversion implemented as a lookup table, reducing it to 8 bits per pixel. The patches are average-pooled (in case of overlap), reordered into raster scan order, and buffered into the dynamic memory by Write Agent 1. Finally, Read Agent 1 conveys the content of the output buffer to HDMI output.

A schematic block diagram of a calculator stage is depicted in Figure 2 (right). Calculations are performed on vectors in fixed point arithmetics with 16 precision bits except the multiply-and-accumulate (MACC) block that uses 48 bit arithmetics internally. To keep a reasonable dynamic range, the data are scaled between various operations by scale factors that were carefully selected to minimize precision loss on a large set of patches from a collection of natural images. Compared to its floating point counterpart, the fixed point implementation produced negligible quality degradation in all our experiments.

Each calculator performs element-wise soft thresholding and the multiplication of the input data by a matrix of size 64×192 (initial stage, converting the input 64-dimensional Bayer patch into a set of 192 coefficients), 192×192 (middle stage, performing operatons on the coefficients), or 192×128 (final stage, converting the coefficients into a 4:2:2 YCbCr patch with 64 luma dimensions and additional 64 chroma dimensions). This is implemented by using MACC blocks of respective sizes. The pa-

rameters of each calculator stage, including threshold values and matrix coefficients, are stored in a local static memory on the FPGA chip.

Since MACC operations are fully pipelined, they require one clock cycle. The total number of clock cycles it takes a single patch to pass though the chain is given by $64 + 192 \times (T - 2)$, where $T - 2$ denotes the number of middle stages. There are additional overheads of approximately 100 cycles per network layer. Due to high resource utilization, we were able to use clock frequency of 125 MHz only. This results in overall throughput of about 16 1920×1080 frames per second without patch overlap and a 4 layer network.

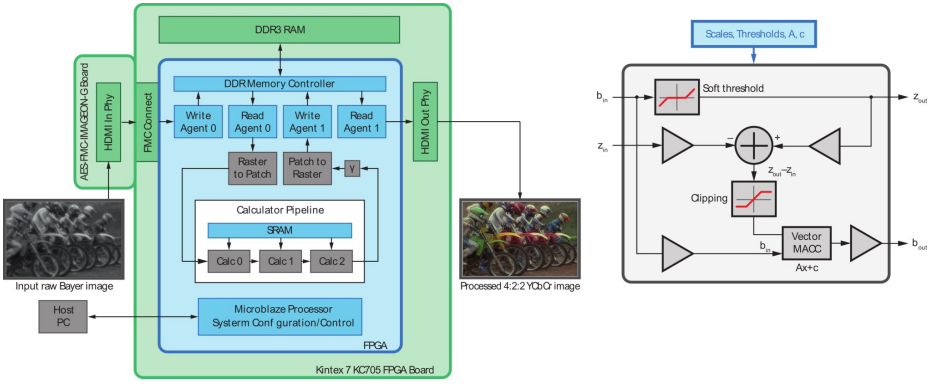

Fig. 2. Schematic description of the FPGA reconstruction system. The raw Bayer image from the sensor at 12bit/pixel is passed, through the HDMI input interface daughter board, to the Kintex 7 FPGA chip. The image is buffered in the external DRAM, from where it is fed as a stream of possibly overlapping 8x8 patches to the calculator pipeline comprising of up to eight stages (see detail on the right), implementing the neural network architecture. The output patches in 4:2:2 YCbCr format are average-pooled and buffered in raster order in the DRAM, from where the image is sent over to the HDMI output interface on the FPGA board. The parameters of the calculator stages and other register values controlling the data flow are stored in the static memory on the chip, into which they are loaded by the host application on system startup.

6 Experimental Evaluation

6.1 Synthetic images

In this experiment, we evaluated the performance of our algorithm on synthetic data of two types. In the first experiment, we used images from the KODAK dataset which were blurred using a PSF simulating typical OOF blur. The same blur kernel was used for the entire image. In the second experiment, we created a synthetic scene with four different depths by using a different kernel in different regions of the image. In all synthetic experiments the reconstruction neural network was trained using 2×10^6 patches taken from the KODAK training set. The network with $T = 8$ layers was converted to fixed-point arithmetics as described in Section 5.

Single Depth. Images taken from the KODAK dataset were used for the evaluation. Each image was convolved with a PSF corresponding to $\psi = 8$ and mosaiced to simulate the input to the system. The algorithms compared were our OMP implementation of [15] with a $k = 192$ atom dictionary trained using k-SVD and our reconstruction neural network with a varying number of layers T. As a reference, we compared our algorithms to the blind deblurring algorithm from [2] following MATLAB default demosaicing algorithm. As the PSF varies across color channels, the algorithm was run separately on each channel. Image reconstruction quality in terms of average PSNR and execution times are presented in Figure 3. It is evident that the highest PSNR is achieved by the neural network with $T = 10$; restricting to $T = 8$ layers performs twice as fast yielding almost the same average PSNR score. Interestingly, the neural network achieves better reconstruction quality compared to the greedy OMP algorithm [15], which we attribute to the effect of supervised training. All sparse prior-based algorithms outperform the blind deconvolution algorithm [2] by over 1 dB PSNR. Comparing the execution times of the algorithms on a standard CPU shows that the OMP algoithm is about 30 times faster than [2], while all neural network implementations are more than 15 times faster than OMP and 500 times fatser than [2] with superior reconstruction quality.

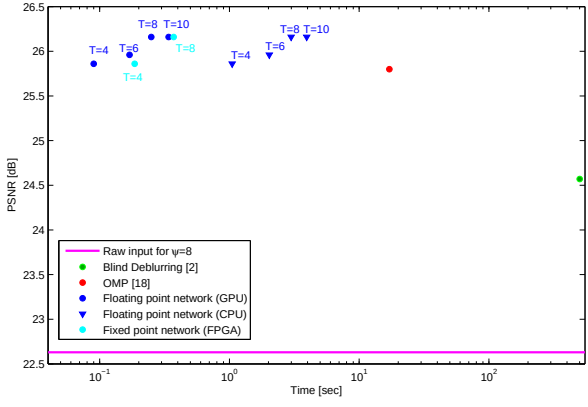

Fig. 3. Comparison of average PSNR and run time on KODAK images. Presented are average values over all test images in the KODAK dataset with simulated blur using $\psi = 8$ (magenta) reconstructed with the different algorithms. Top left corner represents high reconstruction quality and low runtime. One can see that our algorithm (blue and cyan) achieves superior reconstruction quality in comparison to the blind deblurring algorithm of [2] (green) and to the OMP method of [15] (red). In addition, it is evident that increasing the number of layers (denoted by T above the blue and cyan points) improves the performance at an additional computational cost. Our fixed point FPGA implementation (cyan) offers the same reconstruction quality as its floating point (blue) counterpart executed on the CPU while being an order of magnitude faster. Remarkably, the FPGA implementation run-time is in the same range as the GPU floating point implementation. All patch based algorithms were run using a patch stride of 2 pixels. Executing times were measured on an Intel Xeon E5-2650 2GHz CPU and a Tesla C2075 1.15GHz GPU. Our fixed-point implementation was executed on a 100MHz Xilinx Kintex 7 FPGA.

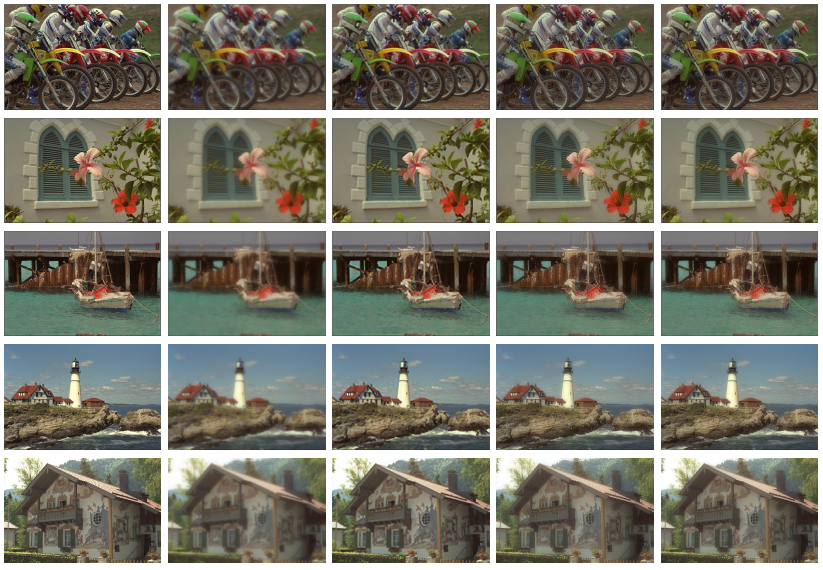

Fig. 4. Synthetic image reconstruction. Columns from left to right: original scene, blurred scene (using a value of $\psi = 8$), blind deblurring using [2], our implementation of the OMP reconstruction of [15] and our FPGA fixed-point neural network with $T = 8$ layers.

In a GPU implementation, our neural network is more than $2,000$ times faster than [2] and 70 times faster than the OMP method of [15]. Furthermore, our fixed point FPGA implementation offers the same reconstruction quality as its floating point counterpart and is an order of magnitude faster than the CPU implementation. Remarkably, the FPGA implementation run-time is in the same range as the floating point GPU implementation, while running at an order of magnitude slower clock and requiring over two orders of magnitude less power. Ground truth images and the reconstruction results by the different algorithms are presented in Figure 4. Qualitatively, one can observe that the fixed-point neural network algorithm with $T = 8$ layers outperforms all other algorithms presented in the figure. Our experiments show that the output of fixed-point network with $T = 4$ is almost imperceavably inferior to the former case, while requiring significantly lower complexity.

Multiple Depths. As indicated in the previous section, for natural depth scenes one cannot assume that the image is blurred by a single blurring kernel. Our reconstruction process analyzes small patches rather than the entire image; therefore, the process is applied to every region inside the image independently of the other regions, allowing our algorithm to treat the input as if it had a single blur kernel. To demonstrate the process, we simulated a "2.5D" scene with four objects each located at a different constant distance from the camera as shown in Figure 5 (top left). The middle image in the top row shows the simulation of the scene acquired using a conventional camera focused on the background. As expected, other objects in the scene are increasingly blurred

with the increase of their distance from the focus point. The rightmost image in the top row presents a simulated acquisition through a phase-coded aperture. The three bottom rows of Figure 5 present zoomed-in snippets of the ground truth compared against reconstruction results using our proposed method, the blind deblurring algorithms [2] and [15] applied to the clear and to the phase-coded aperture images.

Observe that by using our method all objects were restored without noticeable artifacts as opposed to the blind restoration of [2] applied on both aperture types. Quantitatively, our proposed method outperforms all other methods. While the improvement in terms of PSNR over the greedy OMP algorithm proposed by [15] is minor, a noticeable qualitative improvement is apparent in the zoomed-in snippets.

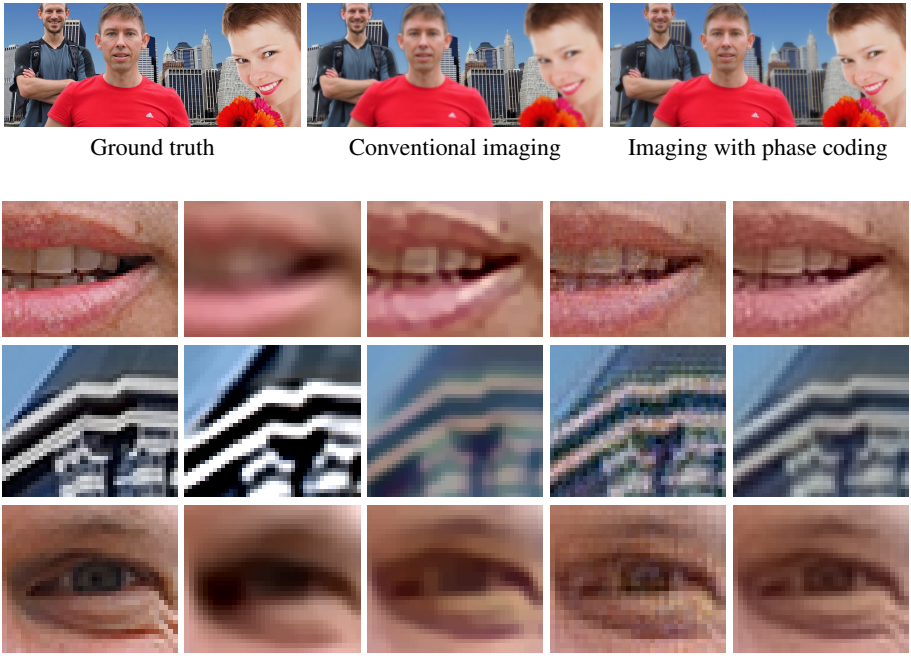

Fig. 5. Synthetic Multiple Depth. The top row shows from left to right: the original synthetic image used for simulating imaging effect on multi depth scene where each object was blurred according to its distance from the camera (the focus point was on the background buildings); Imaging results for conventional clear aperture imaging and imaging with a phase-coded aperture. The three bottom rows present from left to right zoomed-in snippets as follows: Ground truth, blind deblurring proposed by [2] applied on the clear aperture image acquired using a conventional imaging system, blind deblurring of the phase-coded aperture image acquired by our optical system using [2], OMP [15] and our neural network. PSNR values are reported below the images. The full sized images are available in the supplementary material.

6.2 Real optical system

We built a table-top experimental system consisting of a conventional CMOS camera and a 16mm C-mount commercial lens, into which the phase mask was inserted (Figure 6,b).

(a) (b)

Fig. 6. System setup. (a) table setup of the scene presented in Figure 7; (b) 16mm lens with the phase mask (red arrow) inserted into its aperture.

A comparison between a conventional camera and our end-to-end system is presented in Figure 7. The first and second columns from the left show the captured scene with a conventional lens (clear aperture) and with a phase-coded aperture respectively, with the focus set to the background poster. The third and fourth columns show the result of the blind deblurring algorithm from [2] applied to the clear and phase-coded aperture images, respectively. The fifth column presents the scene reconstructed from a phase-coded image using the OMP method from [15]. The rightmost column presents the reconstructed scene using our neural network with 4 layers implemented on an FPGA. Both clear and coded-aperture images were captured in the exact same lighting conditions and exposure time. The entire image regions are shown in the top row whereas the four bottom rows contain zoomed-in snippets of the images at different depths. The superiority of the proposed system is evident in all cases in both quality and run-time. For example, in the magnified fragment of the background (to which the camera was focused), our system produces insignificant changes, while the deblurred clear and coded aperture images show significant over-sharpening artifacts visible as halos around the buildings and cartoon like segments. Our method also outperforms the OMP algorithm proposed by [15] which yields a noisier reconstruction. The bottom row presents a fragment of the resolution chart in which the vanishing contrast spatial frequency was marked with a dashed red line.

7 Conclusions

The phase-coded aperture computational EDOF imaging system presented in this work aims at solving one of the biggest challenges in today's miniature digital cameras, namely, acquisition of images with both high spatial resolution and large depth of field in demanding lighting conditions. Our proposed solution can be easily incorporated to existing imaging systems since it requires the addition of a thin mask (which can be integrated into one of the existing optical surfaces), and a simple real-time hardware

Fig. 7. Full hardware pipeline experiment. From left to right: clear aperture image, phase-coded aperture image, clear aperture image after deblurring using [2], phase-coded aperture after deblurring using [2], phase-coded aperture after deblurring using [15] and the output of our camera using a fixed-point neural network with $T = 4$ layers. The dashed red lines in the bottom row indicate vanishing contrast. The reconstruction time required by [2] was 450 seconds, by the OMP method from [15] 140 seconds while our FPGA only requires 0.5 seconds.

computational unit, which was demonstrated in an FPGA prototype. We showed experimental evidence of the superior image quality of the proposed system compared to conventional acquisition and post-processing techniques. As we demonstrated through extensive experiments, our system outperforms existing techniques for post processing an image taken by a standard camera. It is worthwhile noting that the FPGA system developed for this project is rather flexible and, with moderate adjustments, may be used for implementing other image processing tasks based on sparse prior, such as denoising or restoration from compressed samples.

References

1. Almeida, M.S.C., Almeida, L.B.: Blind and semi-blind deblurring of natural images. IEEE transactions on image processing : a publication of the IEEE Signal Processing Society **19**(1) (jan 2010) 36–52
2. Krishnan, D., Tay, T., Fergus, R.: Blind deconvolution using a normalized sparsity measure. In: Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on, IEEE (2011) 233–240
3. Shan, Q., Jia, J., Agarwala, A.: High-quality motion deblurring from a single image. In: ACM Transactions on Graphics (TOG). Volume 27., ACM (2008) 73
4. Starck, J.L., Pantin, E., Murtagh, F.: Deconvolution in Astronomy: a review. Publications of the Astronomical Society of the Pacific **114**(800) (oct 2002) 1051–1069
5. Couzinie-Devy, F., Mairal, J., Bach, F., Ponce, J.: Dictionary learning for deblurring and digital zoom. arXiv preprint arXiv:1110.0957 (2011)
6. Hu, Z., Huang, J.B., Yang, M.H.: Single image deblurring with adaptive dictionary learning. In: Image Processing (ICIP), 2010 17th IEEE International Conference on, IEEE (2010) 1169–1172
7. Elad, M., Aharon, M.: Image denoising via sparse and redundant representations over learned dictionaries. Image Processing, IEEE Transactions on **15**(12) (2006) 3736–3745
8. Fadili, M.J., Starck, J.L., Murtagh, F.: Inpainting and zooming using sparse representations. Computer Journal **52**(1) (feb 2009) 64–79
9. Dowski, E.R., Cathey, W.T.: Extended depth of field through wave-front coding. Applied Optics **34**(11) (1995) 1859–1866
10. Cossairt, O., Zhou, C., Nayar, S.: Diffusion coded photography for extended depth of field. In: ACM Transactions on Graphics (TOG). Volume 29., ACM (2010) 31
11. Nagahara, H., Kuthirummal, S., Zhou, C., Nayar, S.K.: Flexible depth of field photography. In: Computer Vision–ECCV 2008. Springer (2008) 60–73
12. Cossairt, O., Nayar, S.: Spectral focal sweep: Extended depth of field from chromatic aberrations. In: Computational Photography (ICCP), 2010 IEEE International Conference on, IEEE (2010) 1–8
13. Levin, A., Fergus, R., Durand, F., Freeman, W.T.: Image and depth from a conventional camera with a coded aperture. ACM Transactions on Graphics **26**(3) (2007) 70
14. Chakrabarti, A., Zickler, T.: Depth and deblurring from a spectrally-varying depth-of-field. In: Computer Vision–ECCV 2012. Springer (2012) 648–661
15. Haim, H., Bronstein, A., Marom, E.: Multi-focus imaging using optical phase mask. In: Computational Optical Sensing and Imaging, Optical Society of America (2014) CTh2C–6
16. Goodman, J.W.: Introduction to Fourier Optics. 2nd edn. McGraw-Hill, New York (1996)
17. Milgrom, B., Konforti, N., Golub, M.A., Marom, E.: Pupil coding masks for imaging polychromatic scenes with high resolution and extended depth of field. Optics express **18**(15) (2010) 15569–15584
18. Milgrom, B., Konforti, N., Golub, M.A., Marom, E.: Novel approach for extending the depth of field of barcode decoders by using rgb channels of information. Optics express **18**(16) (2010) 17027–17039
19. Haim, H., Bronstein, A., Marom, E.: Computational multi-focus imaging combining sparse model with color dependent phase mask. Optics express **23**(19) (sep 2015) 24547–56
20. Chen, S.S., Donoho, D.L., Saunders, M.A.: Atomic decomposition by basis pursuit. SIAM journal on scientific computing **20**(1) (1998) 33–61
21. Tropp, J.A., Gilbert, A.C.: Signal recovery from random measurements via orthogonal matching pursuit. Information Theory, IEEE Transactions on **53**(12) (2007) 4655–4666

22. Elad, M.: Sparse and Redundant Representations: From Theory to Applications in Signal and Image Processing. 1st edn. Springer Publishing Company, Incorporated (2010)
23. Aharon, M., Elad, M., Bruckstein, A.: K-svd: An algorithm for designing overcomplete dictionaries for sparse representation. Signal Processing, IEEE Transactions on **54**(11) (2006) 4311–4322
24. Franzen, R.: Kodak lossless true color image suite. see http://r0k. us/graphics/kodak (2004)
25. Brodatz, P.: Textures: a photographic album for artists and designers. Dover Pubns (1966)
26. Beck, A., Teboulle, M.: A fast iterative shrinkage-thresholding algorithm for linear inverse problems. SIAM Journal on Imaging Sciences **2**(1) (2009) 183–202
27. Gregor, K., LeCun, Y.: Learning fast approximations of sparse coding. In: Proc. ICML. (2010)